

Evolution of Things

A.E. Eiben¹, N. Ferreira¹, M. Schut¹, S. Kernbach²

¹Vrije Universiteit Amsterdam, {gusz, nivea, schut} @cs.vu.nl,

²Universität Stuttgart, korniesi@informatik.uni-stuttgart.de

Abstract—Evolution is one of the major omnipresent powers in the universe that has been studied for about two centuries. Recent scientific and technical developments make it possible to make the transition from passively understanding to actively mastering evolution. As of today, the only area where human experimenters can design and manipulate evolutionary processes in full is that of Evolutionary Computing, where evolutionary processes are carried out in a digital space, inside computers, in simulation. We argue that in the near future it will be possible to move evolutionary computing outside such imaginary spaces and make it physically embodied. In other words, we envision the “Evolution of Things”, rather than just the evolution of code, leading to a new field of Embodied Artificial Evolution (EAE). The main objective of the present paper is to offer an umbrella term and vision in order to aid the development of this high potential research area. To this end, we introduce the notion of EAE, discuss a few examples and applications, and elaborate on the expected benefits as well as the grand challenges this developing field will have to address.

I. INTRODUCTION

This article is a position paper about what we call embodied artificial evolution (EAE). The essence of our vision can be briefly summarized as follows. Evolutionary computing as we know it today is disembodied, performed in a digital computer space. However, recent advances (e.g., 3D printing, soft robotics, molecular engineering, synthetic biology, combinatorial chemistry, programmable matter, etc.) make it possible to move evolutionary computing out of the digital space and make it embodied. We believe that this leads to a high potential research and application area that offers great opportunities and poses great challenges. However, to realize the vision, very diverse and presently segregated fields need to interact and cross-fertilize each other. This necessitates a unifying view, an umbrella term and vision to catalyze developments in this direction. This is exactly the main objective of this paper.

Evolutionary computing has long been used as a successful approach for solving optimisation, design, and modelling problems [14], [18]. In this context, evolutionary algorithms are traditionally locked in inside a computer. The individuals that make up the evolving population are digital objects, for example bitstrings, LISP expressions, or artificial neural nets, and all evolutionary operators (reproduction, selection, fitness evaluation) are executed inside the computer. Furthermore, evolution is used as an off-line optimisation or design tool. Even if the object to be optimised or designed is a physical one, it is only created after the evolutionary process terminated. Think of Schwefel’s classic example of an optimally

shaped two phase flashing nozzle or a vibration resistant satellite boom [29].

In this paper we argue that it is conceivable to ‘liberate’ evolution, so that it will take place in the physical world, outside the computer, acting on tangible objects.

The general concept of embodied artificial evolution (EAE) as assumed here can be defined by the following properties.

- 1) It involves physical units instead of a just group of virtual individuals in a computer.
- 2) It has real ‘birth’ and ‘death’, where reproduction creates new (physical) objects, and survivor selection effectively eliminates them.
- 3) Reproduction and selection are not executed through a centrally orchestrated main loop, but in a fully asynchronous and autonomous manner by the individuals themselves. Consequently, the population size may increase or decrease by itself.
- 4) Evolution can be driven by a combination of task-based and open-ended, environmental fitness.

II. DISCUSSION OF THE MAIN NOTION

To aid further elaboration about EAE systems, we consider a number of concrete examples/tasks and use these to illuminate some important aspects of EAE systems.

- 1) The (evolutionary) design of a robot controller for a given robot body and some task(s) in a certain environment.

Here, the objects to be evolved are digital, but are inherently part of a (mechatronic) physical entity. To solve this design problem one could port all evolutionary operators to the robot and execute on-the-fly evolution of controllers. Birth and death, i.e., reproduction and survivor selection, is restricted to the digital space of all possible controllers, on the robot’s processors. However, fitness evaluation happens *in vivo* here as the reproductive probabilities of any given controller are determined by the real-world performance of the robot driven by that controller. In the present literature this approach is called embodied evolution [56] or on-line, on-board evolution [17].

- 2) The (evolutionary) design of a robot body for some task(s) in a certain environment.¹

Here, the objects to be evolved are physical. Thus, one

¹For the sake of simplicity, let us disregard the design of the corresponding robot controller.

could solve this problem by truly embodied evolution, with physical birth and death. In such a system all evolutionary operators work *in vivo*, including reproduction that creates new robots and survivor selection that effectively eliminates them. The main challenge here is obviously formed by the reproduction operators crossover and mutation: how to engineer a system where robots can be born (and die)?

- 3) The (evolutionary) design of a bacterium for some medical or chemical task(s) in a certain environment. Here again, the objects to be evolved are physical. However, while (re)production of mechatronic bodies is a huge challenge, bacteria reproduce by themselves. Thus, that part of the evolutionary machinery is for free in this context. The challenge here is to implement fitness evaluation and the selection operators suited to the given application objectives. Furthermore, one could implement special reproduction operators (mutation and/or crossover) that do not exist in nature, but are useful to solve the given problem.

We can note a couple of things about these examples that help understand some essential aspects of EAE. To begin with, observe that Example 1 is different from Examples 2 and 3 in that it is not truly embodied. To be specific, Examples 2 and 3 illustrate applications where the objects to be evolved are physical. In contrast, the objects to be evolved in Example 1 are digital, only embodied in the sense that they are hosted by a physical robot. Ironically, the term embodied evolution has been introduced for systems like the one in Example 1, cf. [56]. If needed, we can make a distinction by calling this type of systems weakly embodied and using the term strongly embodied for the ones in Examples 2 and 3.

Furthermore, let us note that in case of a robotic application it is possible to separate the body, i.e., the physical robot with its wheels, sensors, etc. and the mind, i.e., the controller regulating the behavior of the robot. Consequently, the task of designing them also can be split into two (and combined, if needed). For the task of designing bacteria, this is not possible, because the regulatory and control mechanisms in bio-chemical organisms are not separated so clearly from the bodies to be regulated.

Yet another difference between a robotic application and a bio-chemical one is the fact that a robotic object is more controllable for the experimenter. Robot bodies are built and robot controllers are programmed by the human experimenters. Even if we consider evolutionary development of robot bodies and controllers, the process is driven by human designed operators. These operators are usually simple; complexity emerges by their interactions. This is not the case for bio-chemical organisms, where the operators are those invented by nature. These are often very complex to understand and to manipulate. For instance, replacing one mutation operator by another one can be easy in an evolutionary robotics application, but switching off one molecular interaction and switching on another one in a cell can be (nearly) impossible.

III. MOTIVATIONS, EXPECTED BENEFITS

There are multiple reasons to investigate EAE systems. First, EAE can lead to solving new design and engineering problems, and solving existing ones in new ways. In fact, EAE technology can be the basis of a paradigm change in how design tasks are solved. Traditionally, the design process of some artifact ends with manufacturing it. Using embodied artificial evolution, design and manufacturing become an intertwined, continuous, on-line activity, propelled by the evolutionary operators (see Figure 1).

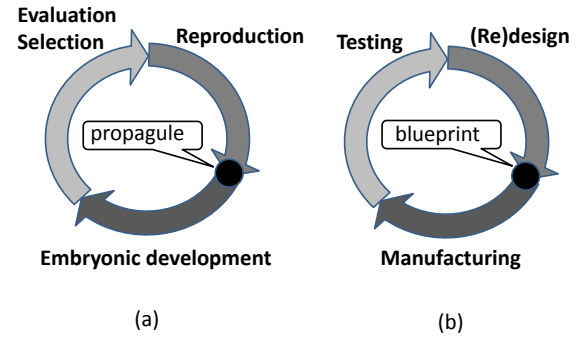


Fig. 1. Two circles showing the analogies between the biological circle of reproduction (a) and the new kind of *in vivo* evolutionary design (b). The effective lifetime is captured by the light gray arrow labeled “Evaluation, Selection” and “Testing”, respectively.

Second, there is much evidence in traditional evolutionary computing showing that evolution can solve problems not solvable otherwise and that evolution can generate unexpected solutions. (Which, then, can be analysed and reverse-engineered, and thus lead to new insights and better understanding of the problem.) Once we equip certain groups of artifacts with the ability to evolve, we create the possibility that some of the evolved designs be truly original, stepping out of the box w.r.t. human thinking.

Third, EAE systems can form the basis of a new experimentalism in biology, where evolution can be studied in a radically new way, based on controlled and repeatable experiments in a new medium. This will enable a deeper understanding of evolution in general, not restricted to (by) evolution-as-we-know-it within the existing life on Earth.

Finally, EAE systems represent an interesting intellectual challenge. The science/art of designing and analysing evolutionary algorithms needs to be reinvented, once we change the medium from purely digital to embodied, physical. For example, population size management is trivial in a genetic algorithm, but preventing an evolving population of robots or bacteria from extinction as well as from explosion can be a hard nut to crack [57]. Furthermore, EAEs mean a great paradigm shift from evolving digital code to *evolving things* in the real world. This implies that the environment where evolution takes place becomes orders of magnitude more complex with inherent randomness (cf. “the noise is

for free”) and a dynamics never encountered in traditional evolutionary computation.

In order to realise these and other unforeseen potential benefits, a lot needs to be done. In the remainder of this paper we elaborate on various related issues. In Section IV we provide an overview of the main areas of research on embodied artificial evolution and in Section V we present three potential applications of EAE technology. In Section VI we elaborate on some of the grand challenges, goals that might serve as catalysts for necessary developments and as stimulus for the specification of research agendas and milestones. Finally, in Section VII we give our concluding remarks.

IV. RELEVANT RESEARCH AREAS

As stated previously, the realisation of embodied artificial evolution systems depends on a number of elements that are still missing. In this section, we describe different research areas that we consider relevant for EAE. For each area, we describe the current state-of-the-art research. We distinguish hard (i. e., mechatrono-robotic) and wet (mainly chemistry and biology) research areas. In fact, systems can be mechatronic, bio-synthetic, bottom-up chemical, or even a hybrid of these technologies. Generally, hierarchical top-down and bottom up approaches inspired by system biology and chemistry are clearly observable; we encounter also multiple examples, such as tissue engineering, neuroscience or reconfigurable hardware, whose complexity represents a horizontal phenomenon from the viewpoint of hierarchies.

A. Hardware: Mechatrono-Robotic Systems

We understand *mechatronics* as the synergistic combination of mechanical engineering, electronic control and information technology in order to design and manufacture useful devices. “It relates to the design of systems, devices and products aimed at achieving an optimal balance between basic mechanical structure and its overall control.” [1].

According to Isermann, an intelligent mechatronic system can be developed and have the ability to model, reason and learn a process (and its functions) in order to achieve a certain goal within a given (time) frame. “An intelligent mechatronic system adapts the controller to the mostly nonlinear behavior (adaptation), and stores its controller parameters in dependence on the position and load (learning), supervises all relevant elements, and performs a fault diagnosis (supervision) to request maintenance or, if a failure occurs, to request a fail safe action (decisions on actions).” [28].

When addressing the challenges associated with the design of mechatronic systems, Alvarez Cabrera *et al.* [8] state that “The challenges are mostly related to integration of design and analysis tools, and automation of current design practices. Addressing these challenges enables the adoption of a concurrent development approach in which the synergetic effects that characterise mechatronic systems are taken into account during design.”

Applications of mechatronic systems, primarily in robotics, strongly involve the concept of embodiment [42],

i. e., using specific properties of materials to achieve a desired functionality. Examples are the locomotion for small jumping robots [32] or embodied sensor-actor coupling [30]. More generally, modern robotics utilises different fields of material science, e.g. [24], which vary from modifications of surface properties up to composite materials with specific mechanical features. We have to mention here also a further miniaturisation of micro-systems [38] and structuring of material by micro-/nano- manipulation [20].

In the literature various references can be found to work related to EAE in computational (mechatrono-robotic) systems. Watson *et al.* in [56], [21] envisioned embodied evolution, somewhat like we do: a population of individuals (in this case, robots) evolves in a completely autonomous manner. Schut *et al* [50] present a somewhat related concept called *situated evolution*, where reproduction creates new minds that become active in a pre-existing robot body, replacing an old one. In [54], Usui and Arita address embodied evolution as in Watson *et al.*: robots evolve based on interactions with the environment and other robots. However, besides transmitting genes when robot encounters take place, each robot executes a genetic algorithm in order to evolve a population of “virtual individuals”. This island model intends to speed up evolution, reducing the influence the number of real robots and the frequency in which robots encounter one another. Nakai and Arita [37] extend this framework by introducing a pre-evaluation mechanism, intended to restrain robot behaviours that are estimated to have a low fitness contribution. Following then same argumentation, Elfving *et al.* in [19] also make use of a subpopulation of virtual agents for each (physical) robot in order to overcome the restriction on population size.

The following examples take the definition of EAE a step further than the previous ones: evolutionary operators and principles are not to be applied to robot controllers only but to the robots themselves. Firstly, Pollack *et al.* present in [43] systems working towards a fully automated manufacturing of autonomous robots. Co-evolution of mind and body of robots is their central issue. The most advanced type of robots are modular robots, where a set of rules that generate the robot structure is evolved. Secondly, another example of modular robot systems is presented in [60], where Zykov *et al.* discuss self-replication: “composed from robotic modules that can change their shape and topological connectivity”, making the overall morphology of a robot adaptable. The replication process starts when a parent creates a new unit, which in its own turn can also assume the role of parent and create other copies of itself.

B. Wetware: Bio-Chemical Systems

Biological systems have a clear advantage over mechatronic devices, as biological properties, such as reproduction, can be taken for granted. For instance, a biological system is naturally equipped to carry out evolutionary processes. Not only reproduction and self-preservation, but also selection and adaptation capabilities are inherent to the system. However, an important challenge is how one can manipulate the

system in order to obtain exactly what one is looking for. Or, in other words, how to *program* a bio-synthetic system? In Section IV-B1 we give some examples on how this has been accomplished so far.

The bio-synthetic systems utilise existing biological cellular systems with their very complex metabolism. The approach from bottom-up chemistry uses another methodology: creating elementary basic cellular (so-called vesicles) and multi-cellular structures “from scratch”. Advantages of this approach are multiple degrees of freedom in designing metabolic networks (in simple cases – autocatalytic reactions) and different internal and external interaction mechanisms. Example of such chemical systems are presented in Section IV-B2.

1) *Top-down Biological Systems*: Programming cells does not have the purpose of substituting silicon computing, but being able to have access to the numerous functionalities and properties on those cells in a predictable, reliable way. Obtaining such programmable bio-devices, or embodied information processors, is a major issue. Potential applications include three-dimensional tissue engineering, biosensing and biomaterial fabrication.

Natural processes can be often described in terms of a networks of simple computational components, or *biobricks* [4]. When referring to biological computing, the main objective is to use the power of natural processes for the purpose of computation. For instance, attempting to “reprogram” parts of a bacterial cell in order to make it perform a certain task. Because natural processes are intrinsically random, changing functionalities of a cell, as well as adding new desired behaviours, is not a trivial exercise. Differently from mechatronic systems, where circuits are built with an specific *known* purpose, the circuitry of a cell may not be completely understood.

However, advances in the area of synthetic biology have allowed some interesting recent results. For instance, in [53], Tamsir *et al* show how logic gates can be built in *Escherichia coli* cells and how complex computations can be produced by “rewiring” communication between cells. “Analogous to a series of electrical gates arrayed on a circuit board, compartmentalisation of genetic gates in individual cells allows them to be added, removed or replaced simply by changing the spatial arrangement of the *E. coli* strains.” [53].

Using an alternative approach, Rigot *et al* describe in [44] how to implement complex Boolean logic computations, which reduces wiring constraints. This is obtained through a redundant distribution of the desired output among the engineered cells. In detail, the engineered yeast cell sense its surroundings (following certain criteria) and send signals to other cells (via secreted wiring molecules). Following the idea of biobricks, a number of yeast cells can then be combined into more complex circuits.

2) *Bottom-up Chemical Systems*: Examples of bottom-up chemical systems can be found in artificial chemistries [15], self-replicating systems [27], using bio-chemical mecha-

nisms for, for example, cognition [12]. In several works, this approach is denoted as swarm chemistry [48]. Researchers hope that such systems will give answers to questions related to developmental models [5], chemical computation [6], self-assembly, self-replication, and simple chemistry-based ecologies [7] or technological capabilities of creating large-scale functional patterns [59]. Several approaches consider meso- and nano-objects, such as particles with functionalised surfaces [49], colloidal systems [23], or molecular networks [39]; a system of elementary autonomous agents, which possess rudimentary capabilities of sensing and actuation. Information processing and collective actuation is performed collectively as, for example, stochastic behavioural rules. Several phenomena, such as meso-scale self-assembling or diverse self-organising processes [13], make these type of systems attractive in applications. Projects such as ECCell, BACTOCOM, NEUNEU, MATCHIT are addressing the questions of chemo-ICT interfaces.

Molecular, colloidal and particle systems also use local interactions and horizontal mechanisms, similarly to 2D and 3D ecological swarms, however we can observe another approach for designing collective phenomena, using the same very simple but large-scale interaction patterns for whole systems [33]. Many research projects in collective nanorobotics and molecular systems are focusing on the technological capabilities of creating such large-scale patterns, e.g., [59].

C. Hybrid Mechatronic and Biochemical Systems

Examples of hybrid mechatronic and biochemical systems are bio-chemical and microbiological systems: using bacterial cellular mechanisms [58] as sensors, the development of bacterial bio-hybrid materials [47], the molecular synthesis of polymers [41] and biofuels [2], genome engineering [10], and more general fields and challenges of synthetic biology [3]. Example of hybrid technologies are attempts to interact with biological populations by means of technological artifacts, for example to manage the grazing of cattle over large areas [51], [11], to control mixed societies of robot and insects [9], or to encourage social communication between robots and chickens [25]. A similar approach is related to the integration of different robot technologies into human societies, for example the management of urban hygiene based on a network of autonomous and cooperating robots [35].

V. APPLICATIONS

The proof of the pudding is the eating. A new technology is largely justified by useful applications. In the present embryonic stage of the EAE field, it is impossible to predict what the best applications will be. However, it is possible to identify some application domains that can demonstrate the benefits of EAE systems in the short/mid-term.

In general, the problems addressed by EAE systems should have a high level of difficulty, including (but not restricted to): (i) changing environments; (ii) multitasking and multi-objective applications; (iii) problems where robust solutions are required; and, (iv) on designing systems with emergent

(self-organising) behaviour. Furthermore, empowered by evolution, systems might present other interesting characteristics such as the capability of learning about and adapting to the environment. Additionally, depending on the application at hand, features such as personalisation of its components, where the machinery becomes adequate to one's personal needs and requirements, may also be of interest.

A. Robot Companions

One could imagine a whole ecosystem of robot companions or artificial pets. Regarding their bodies, these could range from a few cubic decimeters (cat and dog size, if you wish) up to human comparable sizes. As for their mental features, they should be gentle, caring, helpful, and to some extent even intelligent. From a functional point of view they could perform specific tasks, (simple domestic tasks, health monitoring, alarm) and/or provide more generic 'emotional' services (keeping company, being good listeners, acting as partners in simple conversations). Embodied artificial evolution can be a key enabler here, where breeding is the obvious natural analogy. (Note, that breeding as we know it for thousands of years, can be seen a special case of EAE, where mating selection is artificial, while survivor selection and reproduction are natural operators.) As indicated by the examples in Section II, design and implementation of reproduction operators forms a major technical challenge here. Material science is a key discipline to this end.

B. Environment-friendly Organisms

In April 2010 the largest oil spill in US history happened: the equivalent of around 4 million barrels of oil flowed into the Gulf of Mexico, with numerous ecological implications. Analysis on the site, a couple of weeks after the disaster, showed that many groups of bacteria were helping to clean up the waters. These bacteria were able to break down the chemicals found in crude oil and, in fact, responded quite effectively to the incident. However, given the scale of the spill and the fact that some components of the oil could not be broken down by these organisms, the performance of the biological ocean crew was limited. One can imagine a scenario in which such bacteria are synthetically designed. This could have two clear advantages: (i) a large number of organisms could then be generated and deployed, in order to correspond to the scale of the disaster; and, (ii) the bacteria's digestive capabilities could be adjusted in order to guarantee that all oil components could be consumed. In a matter of hours, instead of weeks or months, the waters would be clean, preventing any mid- or long-term consequences for the environment. Such artificially developed organisms can also be used to create building material and biofuel, data storage and stopping desertification.

C. Evolutionary 3-D Printing

Imagine a world in which anything can be produced with just a few clicks. Customised products are at the reach of your hand, ranging from a child's toy to a meal. Vilbrandt *et al.* introduce in [55] the *universal desktop fabrication*

(UDF) that can produce essentially any complete, finished, and functional object. The Fab@Home (www.fabathome.org) is a desktop rapid prototyper and a first step towards UDF. Such personal fabricators can build objects from different materials and will allow anyone to build functional objects that fit its own needs and desires. "You can imagine a 3-D printer making homemade apple pie without the need for farming the apples, fertilising, transporting, refrigerating, packaging, fabricating, cooking, serving and the need for all of the materials in these processes like cars, trucks, pans, coolers, etc."². Additionally, you can think of future fabricators that are able to check what is being printed and adjust the design and mixture of materials in order to obtain the best results. (Embodied) Evolution is expected to play an important role in the development of such fabricators, cf. [45], [46]: "Ultimately, the evolution of form and formation become fully intertwined when the language of assembly itself becomes subject to evolution [...]. Through this co-evolution of form and formation, Evolutionary Fabrication discovers both how to build objects and what to build them out of."

VI. GRAND CHALLENGES

At this moment it is impossible to foresee how this field will develop. However, even in the present early stage, we are able to identify some of the grand challenges that certainly will have to be addressed. In this section we discuss five of these.

A. Body types

The essence of embodied evolution is – the body. To this end, we can distinguish mechatrono-robotic systems (hardware) and bio-chemical systems (wetware), that may also be hybridised. Regarding wetware, there are two options again: bottom-up, relying on chemistry, or top-down, based on biology. Recent developments in microfluidics, functional fluids, or programmable matter also seem very promising. The first grand challenge is thus to find body types suited for (self-) reproduction. Here we can identify several fundamental questions, the most important of those being: *How can the current ICT be combined with bio-chemical developments?* This question is also known in other formulations, as e.g. "programmability of synthetic systems", or "open-ended embodied evolution", and is one of the key points in understanding principles of synthetic life. It is also addressed by the European bio-ICT initiative and several research projects, e.g. PACE [40] and e-FLUX [16], just to name a couple.

B. How to Start – Robot Reproduction

The implementation of birth (reproduction operators) for human engineered physical devices is a critical prerequisite for EAE. These operators must also realise some form of inheritance. The three approaches based on mechatronics,

²Homaro Cantu states in the BBC News article "The printed future of Christmas dinner": <http://www.bbc.co.uk/news/technology-12069495>

chemistry, or biology differ greatly in this respect. (Self-)reproducing mechatronical and chemical units are far from being trivial, whereas it comes for free in biological systems.

In current mechatrono-robotic systems there are two concepts that are crucial for EAE: self-assembling and self-replication. Self-assembling is a process which creates complex systems from basic elements, whereas self-replication means a reproduction of these basic elements. The self-assembling has already been targeted several times in robotics, e.g. [36], [31]. The field of modular (reconfigurable) robotics aims at self-assembling of robot modules into complex structures, so called artificial organisms [34]. Robots are able to make functional copies of artificial organisms from basic building blocks provided there exists an essential reserve of such basic modules.

The self-replication of basic elements remains so far unsolved, even in principle. The problem lies in a high technological complexity of functional units, such as motors, gears or microelectronics. There are several attempts to address self-reproduction in modular robotics, e.g. [60], by additive plastic molding [52] (see also RepRap.org), or by using 3D prototyping technology [46]. No one of these technologies apply self-reproduction as introduced above. For that, further developments are still necessary. This relates not only to the creation of physical units, but also to the computing capabilities, as well as inheritance mechanisms.

C. How to Stop – Kill Switch

A very serious concern regarding EAE is the possibility of ‘runaway evolution’. By this term we do not mean the Fisherian notion of sexual selection reinforcing useless traits [22]. Runaway evolution as we use it here stands for the process of uncontrolled population growth. Such a growth might also be accompanied by the emergence of new, unwanted features in the population. Obviously, it would be highly irresponsible to expose ourselves to such a risk. To reduce this risk, all such experiments could be carried out in highly secured isolated environments, not unlike current research into certain germs, bacteria, viruses, etc. However, this might disable the whole application in cases where the evolving population is inherently free, where the individuals are part of our everyday life (robot companions, waste-eating organisms, medical nano-robots in the human body, etc.). In such cases a ‘kill switch’ is required to guarantee that human supervisors are able to shut down the system, if and when they deem necessary.

D. Evolvability and rate of evolution

It is well-known in biology as well as in evolutionary computing that evolution takes time. One might even say that evolution is slow. Although this statement obviously depends on the context dependent notion of time, it is safe to say that, in general, it takes many generations to achieve a decent level of development. Useful EAE systems must exhibit a high degree of evolvability and a high rate of evolution [26]. In practice, they must achieve decent progress in real time: have short reproduction cycles and/or large

improvements per generation. The main factors here are the application dependent time requirements and the speed of progress, determined by the evolutionary operators. For instance, medical nano-robots put at work in a human body should adapt within a few hours to their environment (the patient’s body). In case of sending evolving robot explorers with a rough initial design to Mars, one can wait months for appropriate designs to emerge.

Building fast evolutionary systems is a nontrivial challenge on its own. Failing to meet this challenge would imply that the real time performance of EAE systems is too low. Ultimately, this could even disqualify the whole approach – at least, for certain applications. In general, the speed of evolution should be used as one of the essential assessment criteria for judging the feasibility of potential applications.

E. Process control & Methodology

A radical change caused by EAE technology is that design and manufacturing become an intertwined, continuous activity. On the one hand, this allows systems to be autonomous and self-improving. On the other hand, this poses an unprecedented challenge for maintaining human control during the process. In particular, human users should be able to perform on-line monitoring and steering in line with the given user preferences. Technically this means directed evolution that could be perhaps realised by directed selection (akin to breeding) and/or directed reproduction (as in genetic manipulation). On a conceptual level, this requires a new kind of methodology that must contain traditional elements, such as specification and validation as well as address previously unforeseen aspects, e.g., mixing (the dynamics of) “free” evolution with specific design objectives on-the-fly. Part of this challenge –or perhaps a challenge on its own– is the ‘freeze switch’, that is, the ability to recognize if/when the evolving objects have obtained the required properties and stop further evolution without killing the system.

VII. FINAL REMARKS

In this paper we have presented the concept of Embodied Artificial Evolution or the Evolution of Things. The systems we envision here are indeed embodied and artificial. They are embodied because evolutionary operators (reproduction, selection, fitness evaluation) are implemented in/by the physical objects that undergo evolution. Furthermore, they are artificial because (i) the individuals and the population as a whole are being designed and/or programmed to fulfill a certain human purpose, to execute a certain task (not excluding open-ended evolution to take place in parallel), and (ii) the evolutionary operators and their particular combination into one working system are human engineered.

We have briefly discussed certain areas of research that can be the basis for developing the physical grounding of such systems. We arranged these areas in a simple taxonomy consisting of hardware (mechatrono-robotic approach), wetware (bio-chemical systems), and their possible hybrid and mentioned a few emerging areas, such as microfluidics, functional fluids, and programmable matter. We have also

provided some potential application examples of EAE and have discussed some of the grand challenges lying ahead, from physical reproduction of robots to monitoring and controlling a running EAE application.

In conclusion, we believe that Embodied Artificial Evolution forms a high potential research and application area. This field is in an embryonic stage, where relevant developments take place within different scientific fields and technological areas that do not naturally interact with each other. We hope that by introducing an umbrella term and a unifying vision we can help bring them together through raising awareness of the shared research issues and possible solutions. As of today, some elements of EAE systems already exist, but considerable scientific and technological advances are necessary to achieve the vision sketched here. However, we do expect that the first examples of such systems will arise in the near future.

ACKNOWLEDGMENTS

We gratefully acknowledge the financial support of the European Commission funding the EvoBody project (grant number FP7-258334), as well as the contribution of the experts we have consulted during the project, cf. <http://www.evobody.eu/>.

REFERENCES

- [1] Mechatronics journal - aims & scope. Elsevier. <http://www.sciencedirect.com/science/journal/09574158>.
- [2] Hal Alper and Gregory Stephanopoulos. Engineering for biofuels: exploiting innate microbial capacity or importing biosynthetic potential? *Nature Reviews Microbiology*, 7(10):715–723, 2009.
- [3] Gil Alterovitz, Taro Muso, and Marco F. Ramoni. The challenges of informatics in synthetic biology: from biomolecular networks to artificial organisms. *Briefings in bioinformatics*, pages bbp054+, November 2009.
- [4] M. Amos. Bacterial computing. In Robert A. Meyers, editor, *Encyclopedia of Complexity and Systems Science*, pages 417–426. Springer New York, 2009.
- [5] J. C. Astor and C. Adami. A developmental model for the evolution of artificial neural networks. *Artif. Life*, 6(3):189–218, 2000.
- [6] Gérard Berry and Gérard Boudol. The chemical abstract machine. In *Selected papers of the Second Workshop on Concurrency and compositionality*, pages 217–248, Essex, UK, 1992. Elsevier Science Publishers Ltd.
- [7] Jens Breyer, Jörg Ackermann, and John McCaskill. Evolving reaction-diffusion ecosystems with self-assembling structures in thin films. *Artif. Life*, 4(1):25–40, 1997.
- [8] A. A. Alvarez Cabrera, M. J. Foekenb, O. A. Tekina, K. Woestenenc, M.S. Erdena, B. De Schuttera, M. J. L. van Tooren, R. Babuškaa, F. J. A. M. van Houtenc, and T. Tomiyamaa. Towards automation of control software: A review of challenges in mechatronic design. *Mechatronics*, 20(8):876–886, December 2010.
- [9] G. Caprari, A. Colot, R. Siegwart, J. Halloy, and J.-L. Deneubourg. Building mixed societies of animals and robots. *IEEE Robotics & Automation Magazine*, 12(2):58–65, 2005.
- [10] Peter A. Carr and George M. Church. Genome engineering. *Nature biotechnology*, 27(12):1151–1162, December 2009.
- [11] Nikolaus Correll, Mac Schager, and Daniela Rus. Social control of herd animals by integration of artificially controlled congeners. In *Proc. of the 10th International Conference on Simulation of Adaptive Behavior (SAB)*. Springer Lecture Notes in Artificial Intelligence LNAI 5040, pages 437–447, Osaka, Japan, July 2008.
- [12] Kyran Dale and Phil Husband. The evolution of reaction-diffusion controllers for minimally cognitive agents. *Artif. Life*, 16(1):1–19, 2010.
- [13] Robert P. W. Davies, Amalia Aggeli, Neville Boden, Tom C. B. McLeish, Irena A. Nyrkova, and Alexander N. Semenov. *Mechanisms and Principles of 1D Self-Assembly of Peptides into β -Sheet Tapes*, volume 35 of *Advances in Chemical Engineering*, pages 11–43. Elsevier, 2009.
- [14] K.A. De Jong. *Evolutionary Computation: A Unified Approach*. The MIT Press, 2006.
- [15] Peter Dittrich, Jens Ziegler, and Wolfgang Banzhaf. Artificial chemistries—a review. *Artif. Life*, 7(3):225–275, 2001.
- [16] e Flux. Evolutionary microfluidix. European Commission, FP7, Future Emerging Technologies project, 2011.
- [17] A. E. Eiben, Evert Haasdijk, and Nicolas Bredeche. Embodied, on-line, on-board evolution for autonomous robotics. In P. Levi and S. Kernbach, editors, *Symbiotic Multi-Robot Organisms: Reliability, Adaptability, Evolution*, chapter 5.2, pages 361–382. Springer, May 2010.
- [18] A. E. Eiben and J.E. Smith. *Introduction to Evolutionary Computing*. Springer-Verlag, London, 2003.
- [19] S. Elfving, E. Uchibe, K. Doya, and H.I. Christensen. Biologically inspired embodied evolution of survival. In Z. Michalewicz and R.G. Reynolds, editors, *Proc. of the 2008 IEEE Congress on Evolutionary Computation*, volume 3, pages 2210–2216, 2008.
- [20] S. Fatikow. *Automated Nanohandling by Microrobots*. Springer-Verlag, London, 2008.
- [21] S. G. Ficici, R. A. Watson, and J. B. Pollack. Embodied evolution: A response to challenges in evolutionary robotics. In J. L. Wyatt and J. Demiris, editors, *Proc. of the Eighth European Workshop on Learning Robots*, pages 14–22, 1999.
- [22] R.A. Fisher. *The Genetical Theory of Natural Selection*. Oxford University Press, Oxford, UK, 1930.
- [23] Masahiro Fujita and Yukio Yamaguchi. Mesoscale modeling for self-organization of colloidal systems. *Current Opinion in Colloid & Interface Science*, June 2009.
- [24] J Gong, L Wan, Q Yuan, C Bai, H Jude, and P Stang. Mesoscopic self-organization of a self-assembled supramolecular rectangle on highly oriented *Proc Natl Acad Sci US A*, Jan 2005.
- [25] Alexey Gribovskiy and Francesco Mondada. Real-Time Audio-Visual Calls Detection System for a Chicken Robot. In *Proceedings of the 4th International Conference on Advanced Robotics*, 2009.
- [26] T. Hu and W. Banzhaf. Evolvability and speed of evolutionary algorithms in light of recent developments in biology. *Journal of Artificial Evolution and Applications*, 2010.
- [27] Tim J. Hutton. The organic builder: A public experiment in artificial chemistries and self-replication. *Artif. Life*, 15(1):21–28, 2009.
- [28] Rolf Isermann. Mechatronic design approach. In Robert H. Bishop, editor, *The Mechatronics Handbook*. CRC Press, 2002.
- [29] A.J. Keane and S. M. Brown. The design of a satellite boom with enhanced vibration performance using genetic algorithm techniques. In I. C. Parmee, editor, *Proceedings of the Conference on Adaptive Computing in Engineering Design and Control*, pages 107–113, 1996.
- [30] S. Kernbach, R. Thenius, O. Kernbach, and T. Schmickl. Re-embodiment of honeybee aggregation behavior in artificial micro-robotic system. *Adaptive Behavior*, 17(3):237–259, 2009.
- [31] Serge Kernbach. Heterogeneous self-assembling based on constraint satisfaction problem. In A. Martinoli and F. Mondada, editors, *10th International Symposium on Distributed Autonomous Robotics Systems*, volume Springer Tracts in Advanced Robotics. Springer, 2011.
- [32] Mirko Kovac, Martin Fuchs, Andr Guignard, Jean-Christophe Zufferey, and Dario Floreano. A miniature 7g jumping robot. In Seth Hutchinson, editor, *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA'2008)*, pages 373 – 378, 2008.
- [33] Sandeep Kumar. Self-organization of disc-like molecules: chemical aspects. *Chemical Society Reviews*, 35(1):83–109, 2006.
- [34] Paul Levi and Serge Kernbach, editors. *Symbiotic Multi-Robot Organisms: Reliability, Adaptability, Evolution*. Springer Verlag, 2010.
- [35] B. Mazzolai, V. Mattoli, C. Laschi, P. Salvini, G. Ferri, G. Ciaravella, and P. Dario. Networked and cooperating robots for urban hygiene: the eu funded dustbot project. In *The 5th International Conference on Ubiquitous Robots and Ambient Intelligence (URAI 2008)*, 2008.
- [36] Shuhei Miyashita, Maik Hadorn, and Peter Eggenberger Hotz. Water floating self-assembling agents. In Ngoc Thanh Nguyen, Adam Grzech, Robert J. Howlett, and Lakhmi C. Jain, editors, *KES-AMSTA*, volume 4496 of *Lecture Notes in Computer Science*, pages 665–674. Springer, 2007.

- [37] J. Nakai and T. Arita. A framework for embodied evolution with pre-evaluation applied to a biped robot. *Artificial Life and Robotics*, 15(2):156–160, 2010.
- [38] B.J. Nelson, L. Dong, and F. Arai. Micro/nanorobotics. In Ous-sama Khatib Bruno Siciliano, editor, *Springer Handbook of Robotics*, pages 411–450. Springer, 2008.
- [39] Jonathan R. Nitschke. Systems chemistry: Molecular networks come of age. *Nature*, 462(7274):736–738, December 2009.
- [40] PACE. *PACE: Programmable Artificial Cell Evolution, FP6*. European Communities, Project reference: 002035, 2004-2008.
- [41] George Pasparakis, Natalio Krasnogor, Leroy Cronin, Benjamin G. Davis, and Cameron Alexander. Controlled polymer synthesis-from biomimicry towards synthetic biology. *Chemical Society Reviews*, 39(1):286–300, 2010.
- [42] Rolf Pfeifer and Josh C. Bongard. *How the Body Shapes the Way We Think: A New View of Intelligence (Bradford Books)*. The MIT Press, 2006.
- [43] J.B. Pollack, H. Lipson, G. Hornby, and P. Funes. Three generations of automatically designed robots. *Artificial Life*, 7(3):215–223, 2001.
- [44] S. Regot, J. Macia, N. Conde, K. Furukawa, J. Kjellén, T. Peeters, S. Hohmann, E. de Nadal, F. Posas, and R. Solé. Distributed biological computation with multicellular engineered networks. *Nature*, December 2010. DOI: 10.1038/nature09679.
- [45] J. Rieffel. *Evolutionary Fabrication: The Co-Evolution of Form and Formation*. PhD thesis, Brandeis University, 2006.
- [46] J. Rieffel and D. Sayles. Evofab: a fully embodied evolutionary fabricator. In *Proc. of the Ninth International Conference on Evolvable Systems*, LNCS 6274, pages 372–380, 2010.
- [47] Eduardo Ruiz-Hitzky, Margarita Darder, Pilar Aranda, and Katsuhiko Ariga. Advances in biomimetic and nanostructured biohybrid materials. *Adv Mater*, 22(3):323–36, 2010.
- [48] Hiroki Sayama. Swarm chemistry. *Artif. Life*, 15(1):105–114, 2009.
- [49] G. Schmid. *Nanoparticles*. Wiley-VCH Verlag, Weinheim, 2004.
- [50] M. C. Schut, E. Haasdijk, and A. E. Eiben. What is situated evolution? In *Proc. of the IEEE Conference on Evolutionary Computation (CEC)*, pages 3277–3284. IEEE Press, 2009.
- [51] M. Schwager, C. Detweiler, I. Vasilescu, D. M. Anderson, and D. Rus. Data-driven identification of group dynamics for motion prediction and control. *Journal of Field Robotics*, 25(6-7):305–324, 2008.
- [52] E. Sells, Z. Smith, S. Bailard, A. Bowyer, and V. Olliver. RepRap: the replicating rapid prototyper-maximizing customizability by breeding the means of production. *Handbook of Research in Mass Customization and Personalization*, 1:568–580, 2009.
- [53] A. Tamsir, J. J. Tabor, and C. A. Voigt. Robust multicellular computing using genetically encoded nor gates and chemical wires. *Nature*, December 2010. DOI: 10.1038/nature09565.
- [54] Y. Usui and T. Arita. Situated and embodied evolution in collective evolutionary robotics. In *Proc. of the 8th Intl Symposium on Artificial Life and Robotics*, pages 212–215, 2003.
- [55] T. Vilbrandt, E. Malone, H. Lipson, and A. Pasko. Universal desktop fabrication. In A. Pasko, V. Adzhiev, and P. Comninou, editors, *Heterogeneous objects modelling and applications*, volume 4889 of *LNCS*, pages 259–284, 2008.
- [56] R. A. Watson, S. G. Ficici, and J. B. Pollack. Embodied evolution: Distributing an evolutionary algorithm in a population of robots. *Robotics and Autonomous Systems*, 39:1–18, 2002.
- [57] W. Wickramasinghe, M. van Steen, and A.E. Eiben. Peer-to-peer evolutionary algorithms with adaptive autonomous selection. In D. Thierens et al., editor, *GECCO '07: Proceedings of the 9th annual conference on Genetic and evolutionary computation*, pages 1460–1467. ACM Press, 2007.
- [58] Janet M. Wood. Osmosensing by Bacteria: Signals and Membrane-Based Sensors. *Microbiol. Mol. Biol. Rev.*, 63(1):230–262, 1999.
- [59] Peng Yin, Harry M. T. Choi, Colby R. Calvert, and Niles A. Pierce. Programming biomolecular self-assembly pathways. *Nature*, 451(7176):318–322, January 2008.
- [60] V. Zykov, E. Mytilinaios, M. Desnoyer, and H. Lipson. Evolved and designed self-reproducing modular robotics. *IEEE Transactions on Robotics*, 23(2):308–319, 2007.